Computational Science at Brookhaven National Laboratory: Three Selected Topics

James W. Davenport, Yuefan Deng, James Glimm, Roman Samulyak

Brookhaven National Laboratory, Upton NY, USA

We present an overview of computational science at Brookhaven National Laboratory (BNL), with selections from three areas: fluids, nanoscience, and biology. The work at BNL in each of these areas is itself very broad, and we select a few topics for presentation within each of them.

1. Fluids

We study hydro- and magnetohydrodynamic (MHD) processes in free surface compressible flows with applications to a variety of Department of Energy projects. Our approach combines theory, numerical algorithm development, simulation based scientific studies, and the analysis of experimental data. Our program has emphasized the study of fluid mixing, and the simulation of unstable interfaces between distinct fluids.

Compressible fluid flows are modeled via the Euler equation. In MHD studies, we are interested in the low magnetic Reynolds number regime, applicable to the study of liquid metals. Negligibly small diffusion time of the magnetic field, typical for such MHD flows, allows the use of a constant external magnetic field approximation [Mol99]. In this case, the governing system of equations can be written as follows

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{u}),$$

$$(1.1) \qquad \rho\left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla\right)\mathbf{u} = -\nabla P + \rho \mathbf{g} + \frac{1}{c}\mathbf{J} \times \mathbf{B},$$

$$\rho\left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla\right)e = -P\nabla \cdot \mathbf{u} + \rho \mathbf{u} \cdot \mathbf{g} + \frac{1}{\sigma}\mathbf{J}^2,$$

where $\rho, \mathbf{u}$, and e are the density, velocity, and the specific internal energy of the fluid, respectively, P is the pressure, $\mathbf{g}$ is the gravitational acceleration, $\mathbf{B}$ is the magnetic field induction, $\mathbf{J}$ is the current density, and σ is the fluid conductivity. The current density distribution can be found using Ohm's law

$$(1.2) \qquad \mathbf{J} = \sigma\left(-\nabla \phi + \frac{1}{c}\mathbf{u} \times \mathbf{B}\right),$$

where the electric potential ϕ satisfies the Poisson equation with the Neumann boundary condition

$$(1.3) \qquad \Delta \phi = \frac{1}{c}\nabla(\mathbf{u} \times \mathbf{B}), \qquad \frac{\partial \phi}{\partial \mathbf{n}} = \frac{1}{c}(\mathbf{u} \times \mathbf{B}) \cdot \mathbf{n}.$$

Here $\mathbf{n}$ denotes a unit normal to the free interface. The viscous and heat conduction terms were neglected in the governing system of equations.

The system of equations (1.1)- (1.3) is a coupled hyperbolic - elliptic system in geometrically complex moving domain. We have developed a numerical method for solving such systems and the corresponding parallel software. The numerical method treats the MHD system in the operator split manner. We use the front tracking hydro code FronTier [ChernGliBry86] with free interface support for solving the hyperbolic subsystem (1.1). The Poisson equation (1.3) is solved using a mixed finite element method on grid dynamically conforming to the moving interface.

FronTier represents interfaces as lower dimensional meshes moving through a volume filling grid [ChernGliBry86]. The traditional volume filling finite difference grid supports smooth solutions located in the region between interfaces. The

location of the discontinuity and the jump in the solution variables are defined on the lower dimensional grid or interface. The dynamics of the interface comes from the mathematical theory of Riemann solutions, which are an idealized solutions of single jump discontinuities for a conservation law. Notice that since we are primarily interested in the contact discontinuity propagation, we restrict ourselves to the Riemann problem for a hydro system of equations and therefore neglect some elementary waves typical for the MHD Riemann problem. Some features of the FronTier hyperbolic solvers include the use of high resolution methods such as MUSCL with a large selection of Riemann solvers. We use realistic models for the equation of state (EOS); some of them were developed to model phase transitions (cavitation) in fluids.

The existence of a tracked surface, across which physical parameters and the fluid solution change discontinuously, has important implications for the solution of an elliptic or parabolic system by finite elements. Strong discontinuities require that the finite element mesh align with the tracked surface in order not to lose the resolution of the parameter/solution. We use a point-shifted technique for the grid generation with interface constraints, a mixed finite element methods based on Whitney elements, and a parallel linear equation solver. The solver is based on the Glowinski-Wheeler domain decomposition, a direct solver for dense linear systems in every subdomain, and a global wire-basket problem for matching solutions across the subdomain boundaries [CowManWhe95]. The method has been implemented as an MHD extension of the FronTier hydro code. In some simple but practically important geometries, the elliptic equation (1.3) may have a simple approximate analytical solution which allows us to skip an expensive step of solving the elliptic problem numerically.

The front tracking code FronTier and its MHD extension have been widely used to study fundamental and applied problems of free surface flows. Numerous numerical studies of the Rayleigh-Taylor and Richtmyer-Meshkov instability problems have been performed. Agreement of the calculated Rayleigh-Taylor mixing rate coefficient with experimental data has been obtained [GeoGliLi02]. Among practical applications of our code, we would like to highlight numerical studies of liquid mercury targets for advanced accelerators.

The design of targets able to generate high-flux beams is among the most important problems of the design of advanced accelerators such as the Spallation Neutron Source (SNS) and the Muon Collider/Neutrino Factory. The need of operating high atomic number materials able to withstand intense thermal shocks has led to the exploration of liquid mercury targets. The proposed Muon Collider/Neutrino Factory target is a free mercury jet interacting with a high energy proton pulse in the presence of a strong magnetic field. Our numerical simulations and recent experiments have shown that strong waves caused by the proton energy deposition lead to surface instabilities and the breakup of the jet [Sam02,SamGliOh03] (Figure 1). We have shown numerically that the magnetic field tends to stabilize the Richtmyer-Meshkov type instability in the mercury jet (see Figure 2) [SamGliOh03].
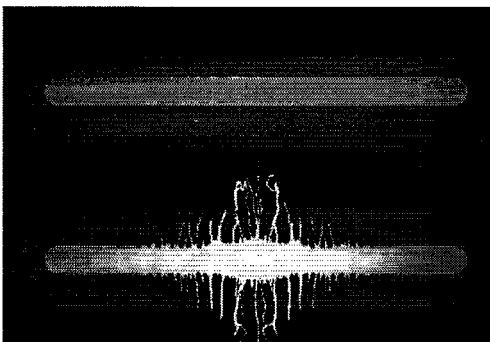

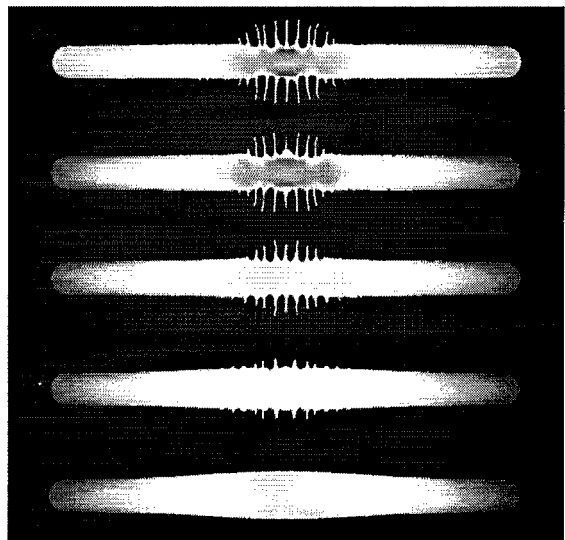Figure 1. Numerical simulation of the Muon Collider target evolution.

Figure 2. Stabilizing effect of the magnetic field on the Richtmyer-Meshkov instability in the mercury jet under the proton energy deposition. a) B = 0, b) B = 2 Tesla, c) B = 4 Tesla, d) B = 6 Tesla, e) B = 10 Tesla.

The numerical simulations shown in Figures 1 and 2 were performed using a single phase equation of state for liquid mercury (stiffened polytropic EOS). Simulations proved that the rarefaction wave strength in the mercury target significantly exceeds the mercury cavitation threshold. Therefore, the formation of cavities takes place in strong rarefaction waves and cavitation bubbles influence the wave dynamics in mercury and the Richtmyer-Meshkov instability of the jet surface. This conclusion was confirmed by experimental measurements of the sound speed in mercury after the interaction with a proton pulse. An accurate description of the wave dynamics in the presence of cavitation depends on the equation of state modeling. It is essential for obtaining quantitative agreements with experiments. We have developed in the isentropic approximation a simple homogeneous two-phase equation of state for modeling cavitation phenomena in fluids. The model is based on the stiffened polytropic EOS for the liquid phase, polytropic EOS for the vapor phase and the following pressure-density relation [Wal69] for modeling the liquid-vapor coexistence domain:

$$P(\rho) = P_l^{sat} + P_{vl} \quad \log \left[\frac{\rho_v a_v^2 \left(\rho_l + \alpha (\rho_v - \rho_l) \right)}{\rho_l \left(\rho_v a_v^2 - \alpha (\rho_v a_v^2 - \rho_l a_l^2) \right)} \right],$$

where ρ_l, a_l, ρ_v, a_v are the density and the speed of sound in the liquid and vapor saturated points, respectively,

$$P_{vl} = \frac{\rho_v a_v^2 \rho_l a_l^2 (\rho_v - \rho_l)}{\rho_v^2 a_v^2 - \rho_l^2 a_l^2}, \quad \text{and} \quad \alpha \text{ is the void fraction:} \quad \alpha = \frac{\rho - \rho_l}{\rho_v - \rho_l}.$$

The corresponding EOS software library was implemented in the FronTier code. The code was applied to the study of the interaction of mercury with an intensive proton pulse in the geometry typical for mercury thimble neutrino factory experiments carried out at the BNL AGS and CERN ISOLDE, and excellent agreement with experimental data was achieved (see Figure 3).
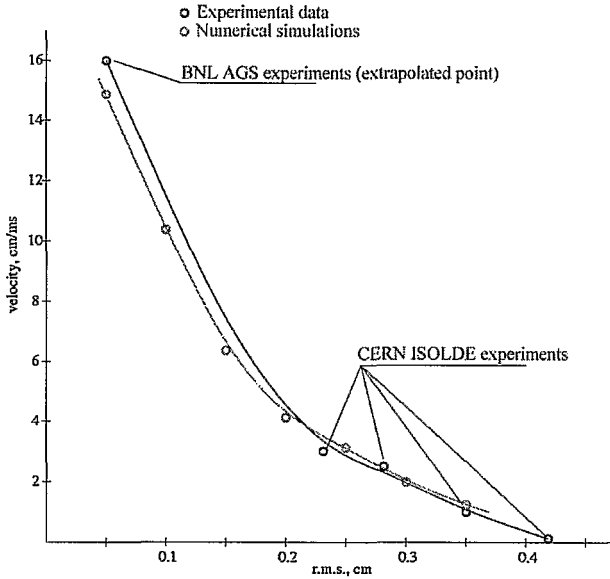

Figure 3. Velocity of the mercury splash in the thimble at 2.5 microseconds after the interaction with the proton beam as a function of the r.m.s. spot size of the beam (beam parameters: 24 GeV, $17*10^{12}$ protons). Blue: experimental data, green: numerical simulations.

A large amount of uncertainty also remains in the design of the SNS target. In addition to a set of difficult engineering problems associated, for instance, with the design of windows able to withstand large thermal gradients and shocks, recent experiments with an SNS target prototype uncovered yet another problem critical to the target lifetime. They showed pitting of stainless steel surfaces that were in contact with mercury subject to large pressure pulses induced by the collapse of cavitation bubbles. This reduces the target lifetime by orders of magnitude. Large-scale numerical simulations of the cavitation in mercury and mitigation techniques capable of reducing the pitting are critical for improvements of the target design and the reduction of the number of costly experiments. Our present studies include the direct (tracked bubbles) approach to cavitating and bubbly flows and the development of an improved full thermodynamic homogeneous equation of state model which use a Rayleigh-Plesset type equation for the evolution of an average bubble and include implicitly the drug, surface tension, and viscous forces, and the mass transfer.

2. Nanoscience

We study nanoscience at the level of Schroedinger's equation for the single particle electron field. The equations are referred to as density functional theory (DFT), widely used for many body quantum mechanics simulations. The central problem in DFT is the solution of the single particle Schrodinger eigenvalue equation

$$(2.1) \qquad H\psi \equiv \left(-\frac{1}{2}\Delta + V \right)\psi = E\psi$$

for the electron wave function ψ, required for the minimization of the Hamiltonian H, or total energy operator. The potential

$$(2.2) \qquad V(\vec{r}) = \sum_R \frac{-Z_R}{\left| \vec{r} - \vec{R} \right|} + \int \frac{\rho(\vec{r}')}{\left| \vec{r} - \vec{r}' \right|} dr' + \mu_K(\rho(\vec{r}))$$

is defined by a sum over the coordinates of the atomic nuclei R with charge Z_R, ρ is the charge density for the electrons, and μ_K is the exchange correlation potential, approximately proportional to $\rho^{1/3}$. In current implementations μ_K also depends on the gradients of ρ [PerBurErn96]. In practice, the electron charge contribution to the potential V, i.e. the term involving the integral of ρ, is evaluated through the solution of a Poisson equation, rather than by direct quadrature, a point we return to later. The two equations are coupled through the formula

$$(2.3) \qquad \rho = \sum_{E_i < E_{\text{Fermi}}} \psi_i^* \psi_i$$

which states that the electronic charge density is given by the sum over the eigenfunctions of H with eigenvalue less than the Fermi energy E_{Fermi}. Thus electron charges ρ not only define H; they are determined directly by it, so that a pair of nonlinearly coupled equations must be solved by an iteration to achieve self-consistency.

A direct implementation of DFT scales as $O(N^3)$ where N is the number of electron states, in practice proportional to the number of atoms. This poor scaling results from the time for eigen analysis of a dense matrix, also with $O(N^3)$ scaling. In addition, scaling of current algorithms is hindered by the use of a delocalizing (plane wave) expansion for the eigenfunctions. Use of local basis elements yields asymptotic decoupling of quantum mechanical effects between remote atoms, reflecting known physical properties of such systems. As the system size increases, we combine this with divide and conquer $O(N)$ DFT scaling, spin dynamics and a coupled spin-continuum approach.

We use the divide and conquer algorithm [Goe99, Yan91, YanLee95], which scales (nearly) as $O(N)$, and is favorable for large systems. In this algorithm, space is divided into disjoint cells, themselves surrounded by a layer (of overlapping) ghost cells. The goal is to determine the charge density ρ correctly in each cell; the separate pieces can then be combined into a global solution of the problem. Starting from a global trial charge density, this is restricted to each cell and its neighboring ghost cells (the extended cell). In the extended cell, the Schroedinger eigenvalue problem is solved. The cell has size $O(1)$ and the cumulative time for this step is $O(N)$. The individual pieces for ρ are glued together, and a Poisson equation is solved to determine a global V. This sequence of operations is repeated until convergence (self consistency) is achieved.

For this algorithm to be successful in practice, we include a number of special steps. First, we use rapidly localized wave functions to define the basis set used to expand ψ (we use Slater type orbitals, defined as powers times exponentials but any localized function will work, including for example Gaussians). Secondly special (overset) grids are used for the solution of the Poisson equation. The space is partitioned into nonoverlapping spheres and an interstitial region. The singular part of V near each atom is solved separately, using a basis defined by spherical harmonics and the radial solution of the Dirac equation. These solutions are matched to a global but smooth and rapidly converging solution based on multigrid methods, using a graded grid to allow open boundaries at infinity.

In order to match the localized basis functions onto the Dirac solutions it is necessary to expand the Slater type orbitals (STO) centered on a given atom about another atomic site. Such expansions are generally called addition theorems. They can be derived using Fourier transforms [Dav84]. The basic idea is to use the Fourier expansion of the localized orbital to perform the translation, then use the Rayleigh expansion of the plane wave.

The STO is given by

$$(2.4) \qquad \phi_{nlm}(\vec{r}) = r^{n-1} \exp(-\zeta r) Y_{lm}(\vec{r})$$

Where Y is a spherical harmonic. Its Fourier transform is obtained analytically in terms of derivatives of spherical Bessel functions times spherical harmonics.

The Rayleigh expansion is

$$(2.5) \qquad \exp(i\vec{q}\cdot\vec{r}) = \sum_{l,m} 4\pi i^l j_l(qr) Y_{lm}^*(\vec{q}) Y_{lm}(\vec{r})$$

Inserting into the Fourier expansion for ϕ leads to

$$(2.6) \qquad \phi(\vec{r} - \vec{R}) = \sum_{L',L''} 4\pi I(L,L',L'') V_{nll'l''}(r,R) Y_{L'}(\vec{r}) Y_{L''}(\vec{R})$$

Here L is a combined index for the pair (l,m) and I is the integral of the product of 3 spherical harmonics (Gaunt integral). These are evaluated efficiently to machine accuracy using Gaussian integration [FerWeiWat94]. The V coefficients are also expressible in terms of powers times decaying exponentials providing the spatial locality which is critical to the favorable scaling of our approach.

With these powerful tools, we will explore new regimes of size with verified fidelity, for novel magnetic properties in dots, tubes (such as transition metal oxides), and surfaces (such as europium sulfide, dots embedded on cobalt). The simulations maintain close coordination with experimental groups.

We are especially interested in the use of DFT to study quantum dots proposed for the design of high density magnetic storage devices. Typical problems involve thousands of atoms with required accuracies of milli or micro electron volts. Thus the divide and conquer algorithm to achieve $O(N)$ scaling and effective parallelism and the use of localized basis functions in the expansion of the Schroedinger equation is important for the efficiency of this approach.

The nanoscience of magnetic dots and tubes and the nanotechnology of the next generation of data storage devices call for quantum calculation of very large systems. Nanoscale patterning offers the promise of dramatically improved densities at lower cost for data storage. Densities of current commercial hard drives are 10^{-5} bits/nm^2 [ThoBes00]. Industry objectives for 2005 call for an increase of a factor of 10 [Mor02]. However, as the dot size decreases, particles become superparamagnetic, losing their ability to retain a magnetic moment in the absence of a magnetic field.

Fig. 4 shows an array of nickel nanodots on cobalt, approximately 5 nm in diameter. If each held 1 bit, the storage density would be increased 40 times to 4 x 10^{-3} bits/nm^2, a density which should be achievable in practical storage devices [Whi00, RosSmiSav99].
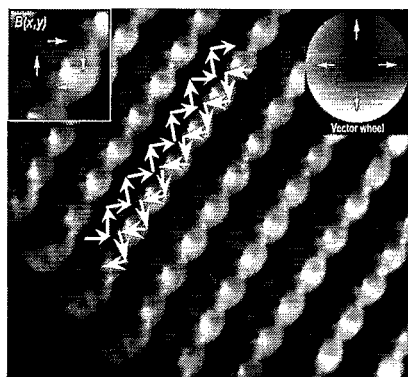
Fig. 4. Magnetic induction of an array of nickel nanodots approximately 5nm in diameter. [Y. Zhu, BNL].

The transition from ferromagnetism to superparamagnetism is determined by the magnetic anisotropy, which pins the magnetic moment in directions aligned to the crystal lattice. The energy differences associated with alignment are extremely small, $\sim 1 - 10$ μeV, so that extremely accurate calculations are needed to distinguish between magnetic and superparamagnetic states, and thus between a dot which is too small and one which is big enough or between a material which is magnetically effective and one which is not. Even the distinction between ferromagnetic and antiferromagnetic states is energetically small (on the order of millielectron volts).

Such distinctions, of basic scientific importance in the study of magnetism; also arise in technology. For example, antiferromagnetic substrate materials are used to enhance the magnetic properties of a ferromagnetic dot. To resolve these differences requires solution accuracy in the range of 1 μeV to 1 meV.

3. Biology

We are interested in structural biology, the docking of proteins and the determination of reaction rates. We use multiple levels of detail in the equations of motion, to allow calculations over time scales of biological interest. The quantum mechanical scale is needed for force calculations where available two body potentials are not sufficiently accurate, as during a reaction event. This will utilize the nano code discussed above. For much of the studies we will use molecular dynamics simulations. These simulations have traditionally been limited in the allowed biological times, due to the inability to achieve fine scale parallelism. We will discuss fine scale parallelism for molecular dynamics and proposed resulting long time simulations (in the multiple microsecond range) based on a novel computer architecture. This is the Quantum Chromodynamics on Chip (QCDOC) architecture designed by physicists for the solution of problems in quantum field theory

We estimate that the QCDOC can integrate molecular dynamics equations for 10^5 particles interacting via long range forces (including Coulomb) for one microsecond of simulated time using less than two weeks of computing time using 8,000 processors. This number of atoms is typical for biological molecules. The two main conclusions we reach are:

1. This is an increase of more than one order of magnitude in simulated time over current simulations.
2. The novel architecture, with 24 parallel channels of low latency communication per processor, allows improved long range communication and an unusual degree of fine scale parallelism, as compared to conventional switch-based architectures.

Here we discuss an analysis of the computing time used in the Ewald method as a function of the required accuracy, the size of the molecular dynamics cell, and the hardware design parameters.

The QCDOC consists of approximately 10,000 IBM PowerPC processors with extremely fast nearest neighbor communication and 4 megabytes of memory located on the each chip with a memory to processor bandwidth of 8 Gb/s It is expected to achieve on the order of 10 teraflops peak (5 teraflops sustained) at a cost of $1 per sustained Mflop. The crucial design feature is a set of 24 serial communication channels to neighboring nodes in a 6D torus communication network, which are capable of sending or receiving data at the rate of 500 Mb/s per channel with a latency of 95 nanosec (send) and 320 nanosec (receive), resulting in an aggregate bandwidth of 12Gb/s for each node. [CheChrCri01].

Now, we consider the QCDOC performance characterization. The computational time $t_{\text{computation}}$ is given by

$$t_{\text{computation}} = N_{\text{ops / processor}} / (e_{\text{efficiency}} \times f)$$ where $N_{\text{ops / processor}}$ is the number of operations per processor, $e_{\text{efficiency}}$ is the single processor efficiency, expressed as a ratio of sustained to peak single processor performance, and f is the single processor peak performance.

Our hardware communication model is likewise very simple. We adopt the model $t_{\text{message}} = b^{-1}x + l$ for the time for a single message sent over a single channel, where $b = 0.5 \times 10^9$ is the bandwidth (bits/sec) per mono-directional channel, $l = 320 \times 10^{-9} = 3.2 \times 10^{-7}$ seconds is the latency and x is the message size in bits.

Let us apply this formula to develop a model for spatially nearest neighbor communication. We need one message hop per coordinate direction to communicate to coordinate-wise nearest neighbors. Thus to reach the off-diagonal coordinate-wise nearest neighbors, we need 3 hops in three spatial dimensions. If $x_{\text{processor}}$ is the per processor message size, then the bandwidth applies to the total of these three messages. There are $27 - 1 = 26$ such units of data to be sent and received at each processor, distributed over the 3 messages. Thus the time t_{nn} for nearest neighbor communication is give by

$$t_{\text{nn}} = 26 \times \frac{2}{wb} \times x_{\text{processor}} + 3l$$ where $w = 12$ is the number of one way wires communicating to coordinate-wise neighbors in the 3D spatial lattice.

Memory management is also an issue. In the interests of brevity, we refer to [DenGliDav03a]

The all to all communication requires further analysis. For an all to all communication, we adopt the same formula, with effective, all to all values for b^{-1} and l,

$$t_{\text{all to all}} = b^{-1}_{\text{all to all}} x + l_{\text{all to all}} \; .$$

The all to all bandwidth is independent of the number of message hops. In fact $b_{\text{all to all}} = wb / 2$ where $w = 24$ is the number of wires per processor ($w / 2$ is the number of wires per direction per processor) and b^{-1} is the single channel one directional bandwidth. Here x is the total message size at a given processor. Thus for P processors, $x = (P - 1)x_{\text{processor}}$.

The all to all latency is $l_{\text{all to all}} = n_{\text{hops}} l$, where l is the single channel latency introduced above and n_{hops} is the number of hops for an all to all message. For a network in the form of a torus in dimension D with each factor periodic (a circle) of size r_i, $[r_i / 2]$ hops are required to traverse a single factor of the torus and

$$n_{\text{hops}} = \sum_{i=1}^{D} [r_i / 2]$$

hops, or $D[r / 2]$ hops for a homogeneous network (all $r_i \equiv r$), to traverse the full network. Here $[a]$ denotes the integer part of the real number a. Thus the time for an all to all communication is

$$t_{\text{all to all}} = \frac{2x}{wb} + n_{\text{hops}} l$$

The success of the QCDOC architecture can be seen from this formula, as the number

$$n_{\text{hops}} = 3\sqrt[6]{P}$$

of hops is small, for P processors arranged in a 6D torus. The total latency $n_{\text{hops}} l$ is small. The communication time is scalable and is dominated by the bandwidth. Here the parallelism of the $w = 24$ fast communication channels per processor becomes important.

In classical MD the positions of a set of interacting particles are advanced according to Newton's second law

$$M_i \frac{d^2 \vec{r}_i}{dt^2} = \vec{F}_i$$ where r_i is the position of the ith particle, M_i its mass and $\vec{F}_i$ the force acting on it due to the other

particles. Here we only discuss the long range forces. The Ewald algorithm [ZhoHarXu01] expresses the Coulomb interaction energy V as

(0.1)
$$V = (1/2)\sum_{i,j}\left(V_{ij}^r + V_{ij}^k\right)$$

where

(0.2)
$$V_{ij}^r = q_i q_j (1-\delta_{ij})\frac{\text{erfc}(\alpha r_{ij})}{r_{ij}}$$

(0.3)
$$V_{ij}^k = q_i q_j \frac{1}{\pi L^3}\sum_{k\neq 0}\frac{4\pi^2}{k^2}\exp\left(-\frac{k^2}{4\alpha^2}\right)\cos(k\cdot r_{ij}) - \delta_{ij}q_i q_j \frac{2\alpha}{\sqrt{\pi}} \, ,$$

The time step optimization results from the use of a smooth cutoff function h to isolate the local and rapidly varying part of the short range forces, through use of the formulas

$$\frac{1}{r} = \frac{1}{r}[\text{erf}(r) + \text{erfc}(r)]$$

(0.4)
$$= \frac{1}{r}[h + \text{erf}(r) + \text{erfc}(r) - h]$$

$$\approx \underbrace{\frac{h}{r}}_{\text{short } \Delta t} + \underbrace{\frac{1}{r}[\text{erf}(r) + (\text{erfc}(r) - h)]}_{\text{Long } \Delta t}$$

Here the first term, h/r, is the local and rapidly varying part of the potential. It is integrated with a fast time step to achieve numerical accuracy. The second term has these same terms explicitly removed.

Computational experiments establish the allowed cutoffs in the finite approximations to the erf and erfc terms above, and the performance model with an assumed single processor efficiency of 30% lead to the projection that 1 microsecond of simulated time can be achieved with femto second time steps for the fast dynamics and 10 femtosecond time steps for the slow dynamics, within an 11 day computation [DenGliDav03].

4. Acknowledgements

5. References

[CCC01] D. Chen, N. H. Christ, C. Cristian, Z. Dong, A. Gara, K. Garg, B. Joo, C. Kim, L. Levkova, X. Liao, R. D. Mawhinney, S. Ohta, and T. Wettig, "QCDOC: A 10-teraflops scale computer for lattice QCD," hep-lat/0011004.

[Dav84] J.W. Davenport. "Linear augmented Slater-type-orbital method for electronic-structure calculations," *Phys. Rev. B* 29, 2896-2904 (1984).

[DenGliDav03] Y. Deng, J. Glimm, J. W. Davenport, X. Cai, and E. Santos. "Performance Models on QCDOC for Molecular Dynamics with Coulomb Potentials," Submitted to *Int. J. High Performance Computing Applications* on January 17, 2003

[DenGliDav03a] Y. Deng, J. Glimm and J. W. Davenport. "Global Communication Schemes on QCDOC" Submitted to *Int. J. High Performance Computing Applications,* 2003

[FerWeiWat94] G.W. Fernando, M. Weinert, R.E. Watson, and J.W. Davenport, "Point Group Symmetries and Gaussian Integration," *J. Comp. Phys.* 112, 282-290 (1994).

[GeoGliLi02] E. George, J. Glimm, X.-L. Li, A. Marchese and Z. L. Xu, "A Comparison of Experimental, Theoretical and Numerical Simulation Rayliegh-Taylor Mixing Rates.' *Proc. Nat. Acad. Sci.* 99, 2587 (2002).

[Goe99] S. Goedecker, "Linear scaling electronic structure methods", *Rev. Mod. Phys.* 71, 1085 (1999).

[Mor02] R. Morris, Storage from Atoms to People, USNIX, Conference on File and Storage Technologies, Monterey CA (Jan, 2002) http://www.usenix.org/publications/library/proceedings/fast02/tech.html

[PerBurErn96] J.P. Perdew, K. Burke, and M. Ernzerhof, "Generalized gradient approximations made simple," *Phys. Rev. Lett.* 77, 3865 (1996).

[RosSmiSav99] C. A. Ross, H. I. Smith, T. Savas, M. Schattenburg, M. Farhoud, M. Hwang, M. Walsh, M. C. Abraham, and R. J. Ram, "Fabrication of patterned media for high density magnetic storage," *J. Vac. Sci. Technol. B* 17, 3168 (1999).

[ThoBes00] D. A. Thompson and J. S. Best, "The future of magnetic data storage technology," *IBM J. Res. Dev.* 44, 311 (2000).

[Whi00] R. L. White, "The physical boundaries to high-density magnetic recording," *J. Magn. Magn. Mater.* 209, 1 (2000).

[Yan91] W. Yang, "Direct Calculation of Electron Density in Density-Functional Theory," Phys. Rev. Lett. 66, 1438 (1991).

[YanLee95] W. Yang and T-S Lee, "A density-matrix divide-and-conquer approach for electronic structure calculations of large molecules," J. Chem. Phys.103, 5674 (1995).

[ZhoHarXu01] R. Zhou, E. Harder, H. Xu, and B. J. Berne, J. Chem Phys., 115, 2348 (2001).